

PDPS General Questions:

- **Won't specifying 'test' in a resource reservation request cause problems if the requester wants to run PGEs when testing?**

Regardless of the type, whether it be 'test' or 'production' or what not, when a resource reservation request is approved and committed, it becomes a ground-event. Then, for the duration of time reserved, the resources will not be available to run any PGEs.

- **Production Requests "PR" Best 1 PR for each day or 3 PRs for each day What is the general rule for entering production requests? Enter them, for example, 1 day per one production request or 3 production requests (1 per shift) per day?**

There is no rule for how many Production Requests (PRs) user can enter through the Production Request editor on a given day. The user determines the number of PRs for each day according to their daily schedule.

- **Orbit file gets updated every ten days. (Production Rules) Who updates this file and how? What is the work flow procedure? [Jeff, Another file that needs updated tracks leap seconds. David]**

It is unknown (at this time) who will have responsibility to update the Orbit file. I would think (since this is an ODL file) that the SSIT operators would get the duty (because they would be familiar with ODL).

- **When entering start/stop times in UTC is this converted to standard and daylight times on time lines or does it stay UTC. Please explain when UTC times are used versus collection times and possible as well wall clock conflict issues?**

Production Request editor uses start/stop times in UTC for collection time and uses LOCAL time for Insertion time (4 PY and Later versions). PREditor internally adjust the daylight saving time so that the operator doesn't have to adjust it.

- **Replanning was not really touched on during training. Can someone provide an overview of Replanning in the current configuration? Then provide an overview of the future Drops. What is the current work around of features that need to be performed for replanning, etc? For example, a current plan is active. A HOT order comes in, what are the procedures required to activate a new plan? The new plan should be no more or less than what is in the old plan, plus the new DPR.**

If there is no previous active plan, then activate all dprs in the new active plan.

If there is a previous "old active" plan:

1. If the dpr is in the old plan but not in the new plan: cancel it. (It will cancel unless the dpr is already running).
2. If the dpr is in both the old and new plan: don't do anything, leave it there.
3. If the dpr is not in the old plan but in the new plan: activate it.

So to add a new dpr "over" the current active plan, create a new plan with all the dprs that are in the current active plan, and ADD the additional dpr. Then activate the new plan.

Note: remember that scheduling occurs at the PR level, not at the DPR level. Therefore, in case 1. if we don't schedule a PR that was in the original plan, *all* dprs in that pr will be cancelled.

- **The entire Resource Definitions part of the Resource Planning job is not very clear. What is the overall strategy for performing this task? Why do you need to define particular component, and in what order do resources needed to be defined? Are there names (proper syntax) that should be used to define various types of resources? I have only seen resources defined during the checkout of PDPS and those values are wide open values. How do I do this for real? For example, defining a dual science processor configuration. Do resources need to be defined for each MODE? Why would one mode be different than another? Should this be a shared function?**

The order of defining resources is as follows:

1. Disks. Although there is no official naming convention, it is wise to use the name of the machine which it is attached to and the mode for which the disk is to be used (OPS, TS1, TS2).
2. Virtual Computers. Again there is no standard naming convention, but you should name them after the real computer or computers upon which they are based and the mode for which they are defined. Then you should associate any Disks that will be used by those Virtual Computers.
3. Real Computers. Again, name them after the "real" computer that they are representing and the mode. You probably want to add something like "rc" to differentiate real computers from virtual computers. Then assign the defined Virtual Computer(s) to the Real Computers.
4. Processing Strings. These can be named any way you want, except the name should probably include the mode. Associate the Virtual Computers with each string.
5. Autosys instance(s).

Disk space and the number of processors per computer DO need to be allocated per mode. The general thinking is that the split should be 80% for OPS and 10% for TS1 and TS2. This can be altered (of course) to suit the needs of the DAAC.

Virtual computer consists of subset of CPUs, RAMs and Disk partitions of a physical computer. This is a way of keeping track of resource use and allocating jobs. If a physical computer (real computer) has a number of virtual computers, the combined total of virtual computers' CPUs, RAMs and disk spaces should make up that of physical computer.

I understand that the 'science processor' is the term used to describe the physical computers at DAAC. If each of the two physical computers contains a virtual computer and if both these virtual computers are linked to a string, and the string to an Autosys, the autosys is configured with dual science processors.

- **What kind of reports can be generated in the current Drop? Specifically, how does the production planner provide the OPS supervisor the number of Production Requests & the number of Data Processing Requests created for a specific day. Not only, how many were created, but how many passed, failed, are in pending status waiting for inputs, or even waiting for science software personnel to evaluate problems? And again, how can the production planner or data management people relate output files that were created have been actually placed in the archive? How can the PGE outputs be related to archive files names?**

OPS supervisor will not be able to generate any reports through planning (PLS) in the current Drop.

- **Why would a production planner plan 1 days worth of work (24 hours) versus 3 shifts (3 x 8=24hr). The actual number of DPRs should be the same. Need to schedule either 1 or 3 PRs for the daily schedule. What makes sense to you? and Why? If this is the case why wouldn't the production planner plan two or three days worth of DPRs in advance? This would seem to make sense when running ASTER data.**

As stated above, this is no hard and fast rule for how many PRs/DPRs should be created at one time. It depends on the type of data being processed. The original design was for 27 hours of DPRs to be planned and activated and at the start of every day.

- **Is there a recovery procedure for each stage of PGE execution within AUTOSYS? For example, what should the operator look for when ALLOCATE fails versus when EXECUTE fails? Has someone defined some of the tips & tricks?**

In most cases the user can attempt to force start the job and it will succeed. See the PDPS Troubleshooting pages

(<http://cmdm.east.hitc.com/baseline/troubleshooting/troubleshooting.html>) for more info on what can cause jobs to fail and how to fix the problems.

- **How does the Production Monitor actually monitor DPR's through the ECS system? It appears that there is not an easy way to tie a pending, waiting, or active DPR to a system status. Obviously a running DPR will be running in AUTOSYS. But, how does the operator know what the status is of a non-active DPR. It could be completed or waiting for input. When waiting for input, how does the operator know what the DPR is actually waiting for? Please define some procedures for this process.**

From Autosys open the OPS Console. On the OPS Console GUI there are two buttons, one that shows Jobs Waiting and one that shows Jobs Completed. These can be used to view jobs that have not yet reached Autosys or have completed and have been taken out of Autosys. There is no easy way to know what a waiting Job is waiting on. The only current way is to go into the database and do a query on the PIDprData table, for all the entries which match the "waiting" DPR, have an IOFlag = 0 and Accepted = 0 (this means that the data has not yet been found).

- **Please explain how monthly, weekly, and daily candidate plans will be used in a production environment.**

You can do candidate plans for any period you might want, but I would strongly advice against activating one for a duration of a month. Let me explain a bit how does replanning work (assuming that it is more than likely that in a month, or a week, you will want to replan your activities).

Lets say you have exploded dprs and made the prs one month long, for example, and the pr contains 30 dprs (one for each day). Now the 15 of the month you decide you don't want to run the production of days 20->30. You have no way of cancelling only from the 20th onwards, you will cancel from the 15th onwards.

- **How can the production planner make a portion of a candidate plan be activated?**

This is kind of an operations question. We don't recommend making such a long plan as a monthly plan. Your candidate plans can be of any length you want, but I don't see it practical to activate plans longer than a day. You will get yourself in a mess of replanning otherwise.

Allow me to explain how replanning works... If plan A is active and then you replan with plan B, the workbench will find all dprs contained in plan A and B for the interval comprising plan A and B (merges the intervals). On a dpr basis it will:

1. If a dpr is both in the old and new active plan, leaves it alone.
2. If a dpr is in the old plan but not in the new plan, it cancels it.
3. If a dpr is not in the old plan but it is in the new plan, it activates it.

If you want to cancel dprs, you will have to make short prs (no longer than a day) and then deal with the scheduling of large number of prs per monthly and weekly plan, and their replanning.

In theory, there is not problem activating thousands of dprs, (except if you overflow the queuing server, depending in the size of it), but we have never tested it and done replanning too.

- **For example, the production planner has a weekly plan and would like to plan just the current days worth of work from the weekly candidate plan and make that the current active daily plan?**

You cannot do that. There is no such concept as partially activating a plan. If the production planner wants to run the current day, then he/she has to create a plan with the activities for that day and activate it.

- **Do not understand how priorities are really going to work. It appears that after PR type, PGE, USER type, and User-select weights are computed to a production priority, the number gets rounded to a whole number. Can someone explain in more detail about what the system default values should be for the various attributes (PR, PGE, USER, User-select) and what which values should be modified for various operation scenarios? Are there actually default values? Has ECS defined these for EDC?**

The operator can update lists of DPR attributes such as PR Type, User Type or PGE type so that each value an attribute can have is tied to a particular priority. Also, the operator can change the weight that each attribute's priority is given. A weight is also given to the priority selected by the user who entered the PR. In this way, a priority can be calculated for each DPR. The final priority of the DPR is used when producing a candidate plan. For the current version, the operator can change the priorities through Production Strategy GUI but WILL NOT get effected when producing a candidate plan

through Planning Workbench. At this time there are no default values for EDC.

- **Did not understand how priorities are really going to work. It seems that priorities are going to be assigned based on Production Request. Therefore, all DPRs within a production request will have the same priority. Can you change the priority of a DPR once it have been placed in the system?**

You can't change the priority of a DPR once it has been placed in the system – after activating the plan.